

Competence and Advice

Anna Denisenko* Catherine Hafer† Dimitri Landa‡

September 10, 2022

DRAFT: Comments Welcome

Abstract

Introduction

Political leaders are inevitably defined by the outcomes of policies they choose, yet those policies are rarely a product of the leaders' *sui generis* judgments. Instead, they are deeply influenced by the actions and judgments of people that the leaders must, through choice or necessity, rely on for advice. It is tempting to assert, in the light of this – as historians and political biographers often do – that the leaders are, thus, only as good as their advisors. But the apparent obviousness of this claim obscures fundamental strategic complexities at the core of the leader-advisor relationship.

We develop a theory of policy advice that focuses on the connection between two central elements of that relationship – advisor competence and the quality of advice that the leader may expect – and calls into question some of the strongly held intuitions about what makes for a good advisor and about the politics of policy-making more broadly. Most significantly, we describe an underappreciated tension between advisors' competence and the informativeness of the advice they may be expected to give and show that, while more competent advisors sometimes give more informative advice, in a wide class of substantively important circumstances we characterize, the right expectation should be the opposite. More broadly, the theory we develop sheds light on when the trade-off between competence and information, or, equivalently, between better and more information, is likely to be more acute and what institutional elements can mitigate or extenuate it.

*Ph.D. student, Wilf Family Department of Politics, NYU, e-mail: ad4205@nyu.edu

†Associate Professor, Wilf Family Department of Politics, NYU, e-mail: catherine.hafer@nyu.edu

‡Professor, Wilf Family Department of Politics, NYU, e-mail: dimitri.landa@nyu.edu

The policy-making setting that is most immediately relevant to our theory is the relationship between political leaders and bureaucrats. In such settings, the leader might be an elected official or a politically appointed head of an established government agency, with agency's career civil servants taking on the role of informed advisor. The following, intentionally somewhat stylized, examples from such settings evoke recent events and help motivate our analysis.

1. The political appointees of a U.S. President are seeking specific information about corrupt government practices in a foreign country X, and the acting U.S. Ambassador to X is in a particularly good position to have such information (or to know how to get it): He is a recognized expert on business and political practices in X and keeps his ear to the ground. Furthermore, because he is understood to be exceptionally capable and a person of integrity, the information he provides would be regarded as reliable and solid. The acting ambassador sees the long-term goal of his work to be supporting democratic institutions and political development in X, which he believes to be central to the U.S. interests in the region and to be best achieved by strengthening the rule of law in X. He also suspects that the political leadership he advises places greater weight on the President's own immediate electoral fortunes than on the success of democratic practices in X, and that the leadership's sudden intense interest in corruption in X is borne of a hope of finding a pressure point to smear a political rival whose family member has done business there. This advisor, then, prefers to "play dumb" and risk dismissal than to reveal to the leader what he knows.
2. The U.S. Consumer Finance Protection Bureau (CFPB) has a staff consisting mainly of civil servants attracted to the mission of "protect[ing] consumers from unfair, deceptive, or abusive practices" from firms in the financial services sector. The CFPB Director is a political appointee, whose preferences reflect the political interests and policy commitments of the sitting president. Suppose that the president, and thus the current Director, are committed to minimizing government regulation of business and to promoting the profitability of big banks. The CFPB civil servants have been crafting specific rules for companies offering consumer credit; they are to offer policy options to the Director, but the Director has the power to choose which option to implement. Suppose the Director wishes to choose the option that will be most advantageous to a subset of banks offering consumer credit, and that the CFPB staff has information on point but does not wish to reveal it because it believes that that should not be a relevant consideration. Instead, it provides information on other features of policy, and claims being unable to obtain the information the leader desires.

3. A Prime Minister of country Y strongly supports restarting the program of targeted assassinations of known leaders of terrorist organizations, which he values, in part, because they give him a chance to look tough and effective in protecting the country. Suppose that Y's head of an intelligence agency, on the other hand, has come to believe that targeted killings of terrorist leaders provide relatively little benefit in disrupting terrorist activity and distract the agency from its mission of developing intelligence about potential terrorist threats. While the Prime Minister presses his head of intelligence agency (the advisor) to identify the location of a desired target, the advisor may prefer to forgo gathering the relevant information or to fail to reveal it, even at the risk of angering the leader.

In these and many other instances, civil servants have preferences that suggest an interest in maintaining status quo. Lawrence J. Peter (of the Peter Principle) puts this in a characteristically flip fashion: “Bureaucracy defends the status quo long past the time when the quo has lost its status.” But while Peter’s comment implies a certain degree of criticisms, Huq and Ginsburg (2018) view this conservatism as a guarantor of political stability, and describe as essentially symmetric with respect to possible challenges to status quo: “Just as bureaucracy may make progressive reform difficult to achieve, it also slows down rapid shifts away from liberal democratic norms” (p. 129). The bureaucracies’ status-quo bias need not imply their primitive conservatism – as a rule, bureaucrats have long time horizons and so are apt to discount the value of responding to what may be short-term trends. In contrast, the time horizons of elected officials tend to be short, and they prefer to match the policy to the immediate trends, pandering to the voters’ prior beliefs or otherwise bolstering their electoral fortunes while discounting longer-term implications of their actions.

In our model, an advisor, who, like the bureaucrats in this description, prefers maintaining the status-quo policy and has access to superior information, chooses whether to reveal to the leader her privately observed signal (i.e., send a “verifiable message”) about the state of the world – a signal relevant to the leader’s choice of policy. Assuming the advisor shares her signal, the leader updates more strongly if the advisor is commonly known to be more competent and, conditional on such an update, shifts policy farther. Given this expectation and the gap in most preferred policies between the leader and her advisor, a more competent advisor stands to lose more from revealing her signal than does a less competent one. In general circumstances, the less competent advisors will, thus, have stronger incentives to reveal their information to the leader than will more competent ones. The leader, then, faces a trade-off between the quality of advice she receives and the likelihood of receiving advice.

We describe conditions that influence the advisors’ incentives to share and obtain in-

formation as well as conditions that determine when leaders may be better off having as their advisors agents with lower competence, even if that means that the advice they receive is less reliable. We also show that the possibility of whistle-blowing, which may allow the leader to obtain the information available to the advisor even when the advisor is seeking to conceal it, and which, on its face, appears to be an effective way to counteract the advisor's incentives to conceal, ultimately has an equivocal effect on leader's utility. Holding fixed advisor's information, whistle-blowing encourages information-sharing by the advisor and enables the leader to reach for higher quality advisors, but the possibility of whistle-blowing also discourages information acquisition by more competent advisors and when it is sufficiently likely, becomes detrimental to the leader's information.

The rest of the paper proceeds as follows. Following a brief review of the prior literature, we develop a general formulation of the trade-off between advisor competence (quality of information available to the advisor) and the strategically supportable quantity of advice. We then study a model with quadratic utilities which, we show, generates this trade-off in equilibrium. After that, allow for endogenous information acquisition and let the Advisor choose both whether to conduct the experiment and the experiment's precision. We show, first, that the only precision consistent with equilibrium play is greatest possible precision, and that the Advisor's competence affects her incentives to pursue information in a way that echoes the incentives to reveal information in the fixed-information environment, but that such incentives are moderated by the nature (private vs. public) of the experiment. In the last part substantive section, we introduce the possibility of whistle-blowing and study its effects on the Advisors' incentives to share and to acquire information.

Connection to the Literature

The relationship between leaders and their advisors is critically affected by three sources of agency problems: (1) advisors have their own policy preferences; (2) they have informational advantage over the leaders – the key reason for leaders' need of their services, but, simultaneously, also an impediment to the leaders' ability to evaluate the advice; and (3) advisors may have differing abilities to obtain information, which affects the quality of advice they could give to the leaders. The primary focus of the previous work has been the relationship between the first two of these three factors – difference between leaders' and advisors' preferences and the advisors' informational advantage (the latter sometimes attributed to differences in incentives for rational ignorance: it would take the elected officials' time and effort to delve into the diverse set of issues that civil servants encounter every day (e.g., Downs 1957; Mizrahi and Minchuk 2019)). An important review of this literature is Sobel

(2013).

Information revelation through advice, including from multiple senders, has been a focus of substantial cheap-talk literature in political economy, including Gilligan and Krehbiel (1989), Austen-Smith (1990), Austen-Smith (1993), and Battaglini (2002). This literature shows that the divergence in the actors' preference limits revelation, and successful communication occurs only when there is little preference heterogeneity. In the context of an interaction between a legislative committee and the less informed median voter on the legislative floor, Gilligan and Krehbiel (1989) show that closed rule, which limits the amendments that the floor can make to the committee's proposal before the up-or-down vote, can increase the committee's incentives to reveal. In contrast with this literature, this paper models the communication of hard evidence, rather than cheap talk, and focuses not on the effects of preference heterogeneity, but, holding the difference of the actors' preferences fixed, on the effects of the advisor's competence on her incentives to reveal and to obtain information. Further, unlike in this literature, the sender's preference is to maintain the status quo, and so against revealing any information to the receiver, making closed rule moot in this setting.

In the verifiable messages (persuasion) games literature on advice to which the current model belongs, an important early paper is Shin (1994), which shows that the unraveling in the event of "no news" fails when the advisor's knowledge is imperfect; see also Wolinsky (2003). Dziuda (2011) studies a setting where the fixed expert's preferences are different and unknown to the decision-maker. She shows that there is never full disclosure, but the expert offers pros and cons for the advocated alternative in order to pool with the honest/non-strategic type. Che and Kartik (2009) investigate a situation in which the decision-maker and the advisor are assumed to have identical preferences, but different priors, and the advisor invests effort into acquiring information. They show that the decision-maker prefers an advisor whose prior beliefs are different than her own to incentivize the advisor to acquire information in order to persuade the decision-maker.

Bhattacharya and Mukherjee (2013) and, extending their model, Bhattacharya, Goltsman, and Mukherjee (2018) study verifiable advice from a panel of experts who may vary in quality and preference. Unlike in our model, the focus in these studies is on the optimal extent of conflict between multiple experts, who are assumed to observe the state directly with some probability (their quality). Bhattacharya and Mukherjee (2013) identify the possibility that improving advisor's quality can lower the decision-maker's utility. A necessary condition for such an outcome in their model is that the leader's default policy in the absence of revelation be sensitive to advisor quality. In contrast, in our model with an unbiased leader, the optimal default policy is constant in advisor quality, and the impact of the latter is, rather, channelled through the quality of information provided, allowing us to get a sharper

characterization of the conditions under which the leader may prefer a lower-quality advisor and more fully examine its implications.

Unlike the above studies, our focus is on the strategic implications of the relationship between the advisors' informational advantage over leaders and the quality of advice they could give to the leaders. A standard intuition, captured in a number of political economy models, is based on the career-concerns rationale: the agent wants to appear well-informed (competent). See Ottaviani and Sørensen (2006) for an explicit analysis of the effects of this motivation. By contrast, in our model, the advisor's incentives are such that being well-informed can make it less likely that the advisor is selected, and, conditional on being selected, is all downside.

Several papers study mechanisms, different from the one we analyze, suggesting a possible downside of the advisor/agent competence. Egorov and Sonin (2011) explore the competence-loyalty trade-off in the relationship between leaders (dictators) and advisors (viziers). In their model, the higher is the Vizier's competence, the more confident he is that his betrayal of the Dictator will lead to the Enemy's victory, and, thus, the costlier it is for the Dictator to enforce the loyalty of more qualified viziers against the Enemy's offer of a bribe to the Vizier. The Dictator, thus, faces a loyalty-competence trade-off with respect to the Vizier. See also Terai and Glazer (2018). Sobel (1993) shows that when the agent is uninformed, he might exert more effort to achieve an outcome than a better informed agent, leading to the possibility of the former being more attractive to the principals. Gailmard and Patty (2007) model bureaucratic competence as an exogenous cost of acquiring expertise and show that when cost is too high, the legislature does not reward the expertise acquisition.

There is a growing body of empirical work on competence-loyalty trade-offs. This includes Abbott et al. (2020), who informally describe a conflict between competence and control faced by the governors: while the governors are assumed to prefer to work with highly competent intermediaries, the more competent the intermediary is, the more likely he is to use policy benefits to free himself of the governor's oversight. Other papers in this literature, focusing especially on Chinese politics, include Bai and Zhou (2019), Reuter and Robertson (2012), Shih, Adolph, and Liu (2012), and Xi (2018).

Finally, our analysis connects to a small political economy literature on whistle-blowing. A broadly accepted intuition behind protections of whistle-blowers is based on the belief that whistle-blowers expose irregularities or fraud from within the organization (ADD CITES). Ting (2008) explores an efficiency-based rationale for such protections: when employee's and the politician's preferences are aligned, employee's whistle-blowing enables implementation of projects that would otherwise be vetoed by a biased advisor. Our model suggests a novel reason behind whistle-blower protections: the existence of whistle-blowing improves

the Leader’s utility premium from having a more competent Advisor.

The General Environment

We analyze a strategic interaction between a Leader (he) and an Advisor (she) of known competence. The Leader wishes to choose an action that will match a state of the world, which he does not directly observe. Instead, the Leader may be able to obtain information about the state from his Advisor, whose competence determines the informativeness of the signal about the state of the world that the Advisor privately observes. The timeline of the game is as follows.

1. Nature determines the state of the world $w \in \mathbf{R}$, where w is a draw from a normal distribution $N(\mu, 1/q)$ parameterized by mean μ and precision q , both of which are known.
2. The Advisor of known competence θ observes signal s^A about the state of the world w , $s^A = w + \varepsilon$. The variable ε represents random noise drawn from a normal distribution with mean 0 and precision θ (s.t. $\theta \in \mathbf{R}^+$), $\varepsilon \sim N(0, 1/\theta)$.
3. The Advisor chooses which message m to send to the Leaders, $m \in \{s^A, ?\}$.
4. The Leader observes message m and decides which policy $a \in \mathbf{R}$ to implement.

We denote the Leader’s and the Advisor’s preferences by $U_L(\cdot)$ and $U_A(\cdot)$ correspondingly. We assume that both the Leader and the Advisor have single-peaked preferences. The Leader wants to match the state of the world,¹ while the Advisor wishes to sustain the status-quo, which is set at 0, the expected value of the state of the world.

Following the Advisor’s message m , the Leader forms posterior belief $\mu_1(w|m, \mu, \theta)$, where $\mu_1(\cdot)$ denotes the probability that the state of the world is w conditional on the message m the Leader observes and the competence of the Advisor θ . The Leader’s utility depends on the state of the world and the policy a he chooses. The Leader chooses $a(m, \theta)$ such that it maximizes his expected utility

$$a(m, \theta) = \arg \max_a \int U_L(a, w) d\mu_1(w|m, \theta). \quad (1)$$

¹We focus on this formulation here in order to cleanly state the trade-off between quality (competence) of advisors and their willingness to reveal their privately held information. Below, we will also consider the specification in which the Leader has a bias in favor of moving a policy in one as opposed to another direction.

The Advisor's utility depends on the policy a that the Leader implements and on whether or not the signal she sends is informative. We assume that every Advisor values office and receives a finite and positive utility Ψ while in office. We assume, further, that if the Advisor does not send the informative signal (i.e., if $m = ?$), the Advisor is immediately replaced. By way of justification, one might imagine that competence is not possessed by a single Advisor but instead by the entire body of public servants available. For instance, it might be that the Advisor is part of a bureaucracy, and regulations and formal rules constrain the information she observes. Therefore, if the Advisor is dismissed, the same constraints will affect her successor. Because of that, it becomes sequentially rational for the Leader to replace the Advisor for failing to provide an informative message or to commit to rewarding the Advisor who does send such a message.

The Advisor gets utility

$$U_A(a, m) = \begin{cases} u_A(a), & \text{if } m = ?, \\ u_A(a) + \Psi, & \text{else.} \end{cases} \quad (2)$$

The Advisor sends message $m(\theta)$ that maximizes

$$m(\theta) = \arg \max_{m \in \{s, ?\}} U_A(a(m, \theta)). \quad (3)$$

Finally, the Leader selects the policy. We deliberately assume that the Leader derives no direct utility from the Advisor's competence; yet, the competence indirectly affects the Leader as it alters signal informativeness.

The solution concept is Perfect Bayesian Equilibrium, requiring satisfaction of conditions (1) and (3), as well as efficient posterior $\mu_1(\cdot)$.

The Trade-off Between Competence and Advice

Impact of Competence on Revelation

When the Leader observes the informative signal ($s \in I^L$), the Leader adopts policy equal to mean of the posterior distribution $a = \frac{q\mu + \theta s}{\theta + q} = \mu + (s - \mu) \frac{\theta}{q + \theta}$. Otherwise, he implements a default policy ($a = d$) in the absence of verifiable information.

The Advisor's incentives to share information with the Leader, given the Leader's anticipated policy response, vary with the Advisor's competence. The informed Advisor's utility

from revealing her information ($m = s$) is

$$u_A(\mu + (s - \mu)\frac{\theta}{q + \theta}) + \Psi.$$

Alternatively, her utility from concealing the information is

$$u_A(d).$$

The Advisor reveals her information if and only if

$$u_A(\mu + (s - \mu)\frac{\theta}{q + \theta}) + \Psi > u_A(d). \quad (4)$$

The Advisor's strategy is a mapping from the signal she observes into decision to reveal information or not. Note that the Advisor's strategy does not depend on her probability of observing information ρ . Because the Advisor wishes for policy to match the status quo, her utility $u_A(x)$ decreases in $|x - \mu|$. She follows a threshold strategy and reveals a signal if and only if it belongs to an interval $(\hat{s}(\theta), \bar{s}(\theta))$, where $\hat{s}(\cdot) = \mu - \bar{s}(\cdot)$.

Lemma 1. *The lowest and highest signals that the Advisor reveals to the Leader, $\hat{s}(\theta)$ and $\bar{s}(\theta)$, are symmetric around the ex ante expected value if the state of the world, μ .*

Importantly, because $u_A(\mu + (s - \mu)\frac{1}{q/\theta + 1})$ decreases in θ for every s such that $s < \mu$ and increases in θ for $s > \mu$, $\hat{s}(\theta)$ increases in θ , and $\bar{s}(\theta)$ decreases in θ . Therefore, $r(\theta) = Pr[s \in (\hat{s}(\theta), \bar{s}(\theta))]$ decreases in θ : more competent Advisors are less likely to reveal the information that they observe.

Proposition 1.

1. *The Advisor's incentives to send an informative message to the Leader decrease in the Advisor's competence (θ) and increase in the precision of the prior (q).*
2. *The Advisor's incentives to send an informative message to the Leader increase in the Advisor's valuation of office (Ψ).*

The intuition behind the first part of Proposition 1 is straightforward. The final policy the Leader adopts is a weighted combination of information he knows (his prior beliefs about the state of the world) and information he learns from his Advisor. The higher the Advisor's competence, the more the Leader relies on the informative message. Therefore, when the signal the Advisor observes differs from her most preferred policy, her incentive to reveal it will decrease with her competence, as an informative message from a more competent

Advisor increases the distance between the final policy and the status-quo policy, which the Advisor prefers. For the same reason, the more precise is the prior, the more likely is the Advisor to reveal her signal to the Leader: the higher is q , the less the Leader updates based on the message he receives.

The second part of Proposition 1 highlights the impact of the the benefit of retaining her office (Ψ) on the Advisor's decision to reveal her signal. The more highly the Advisor values her position, the higher is the opportunity cost of concealing information from the Leader, and, thus, the more likely is the Advisor to send an informative message.

Value of [in]Competence

This section focuses on the expected impact of the Advisor's competence on the Leader's utility. To provide a general intuition for the benefit to the Leader of having a less than maximally competent advisor, we abstract away from the exact functional forms of the utilities and signal distributions. Consistent with the results above, the Advisor sends an informative signal to the Leader with probability $r(\theta)$ and sends an uninformative signal with complementary probability $(1 - r(\theta))$, where $r(\theta)$ depends on Advisor's competence. When the Leader observes the informative signal, he gets the expected utility $\alpha(\theta)$. When the Leader does not, he gets expected utility $\beta(\theta) < \alpha(\theta)$.

The Leader's expected utility is

$$E[U_L(\theta)] = r(\theta) \alpha(\theta) + (1 - r(\theta)) \beta(\theta).$$

Note that under certain conditions, the Leader prefers a less competent Advisor over a more competent one. His expected utility decreases in θ when

$$\frac{\partial E[U_L(\theta)]}{\partial \theta} = r'(\theta) (\alpha(\theta) - \beta(\theta)) + r(\theta) \alpha'(\theta) + (1 - r(\theta)) \beta'(\theta) < 0, \quad (5)$$

and increases in θ otherwise.

The next proposition follows from equation 5:

Proposition 2. *The Leader's utility decreases in the Advisor's competence when*

$$r'(\theta)(\alpha(\theta) - \beta(\theta)) < r(\theta) \alpha'(\theta) - (1 - r(\theta)) \beta'(\theta), \quad (6)$$

and increases in her competence otherwise.

This proposition holds that the Leader is better off with a less competent Advisor than with a more competent one when the marginal utility gained from acquiring information

is smaller than the marginal loss of information per se multiplied by its importance to the Leader.

Of course, the question of whether these conditions are consistent with a properly micro-founded agency model remains open. In the next section, we describe such a model with quadratic utilities and show when such conditions are supported in the context of that model.

Quadratic Utilities Model

We begin with a model in which, like in the general setting described above, the Leader and the Advisor interact prior to the Leader's policy-making. The Advisor learns information that is not directly observable by the Leader and has to decide whether to reveal this information to the Leader. In this section, however, we assume that both the Advisor and the Leader have quadratic utilities, which allows us to deepen our understanding of the conditions necessary to guarantee that, in equilibrium, the Leader's utility is maximized at an interior value of the Advisor's competence. We start with the assumption that the Advisor's preferences are state-independent and later show that the results extend to the advisors with state-dependent, but more conservative than Leader's, preferences. Throughout, for simplicity, we normalize the prior signal distribution to have mean $\mu = 0$.

Incorporating the Leader's assumed replacement strategy, the Advisor gets utility

$$U_A(m|s, \theta) = \begin{cases} (a - 0)^2 & \text{if } m = ?, \\ (a - 0)^2 + \Psi & \text{else.} \end{cases} \quad (7)$$

The Leader gets utility

$$U_L(a|m, \theta) = (a - w)^2. \quad (8)$$

The Leader acts last. He chooses action $a(m) = \frac{m}{1+q/\theta}$ when he observes an informative message $m = s$ and chooses action $a(m = ?) = d(\theta)$ otherwise. The Advisor sends an informative message when she observes a signal s such that

$$\sqrt{\Psi + d(\theta)^2} \left(1 + \frac{q}{\theta}\right) < s < \sqrt{\Psi + d(\theta)^2} \left(1 + \frac{q}{\theta}\right) \hat{s}(\theta, \Psi) \quad (9)$$

and sends no informative message otherwise. As in the generalized setting, the Advisor's incentive to reveal her signal to the Leader decrease with her competence. Note that, in the absence of informative message, the Leader's optimal strategy is to implement policy $a(m = ?) = 0$, as the Advisor's strategy is symmetric around 0.

Figure 1 shows the thresholds $\hat{s}(\theta, \Psi)$ as a function of the Advisor's competence θ . The

shaded area depicts signals that an Advisor of competence θ reveals to the Leader. The higher is the Advisor's competence, the smaller is the range of informative messages the Advisor will send to the Leader.

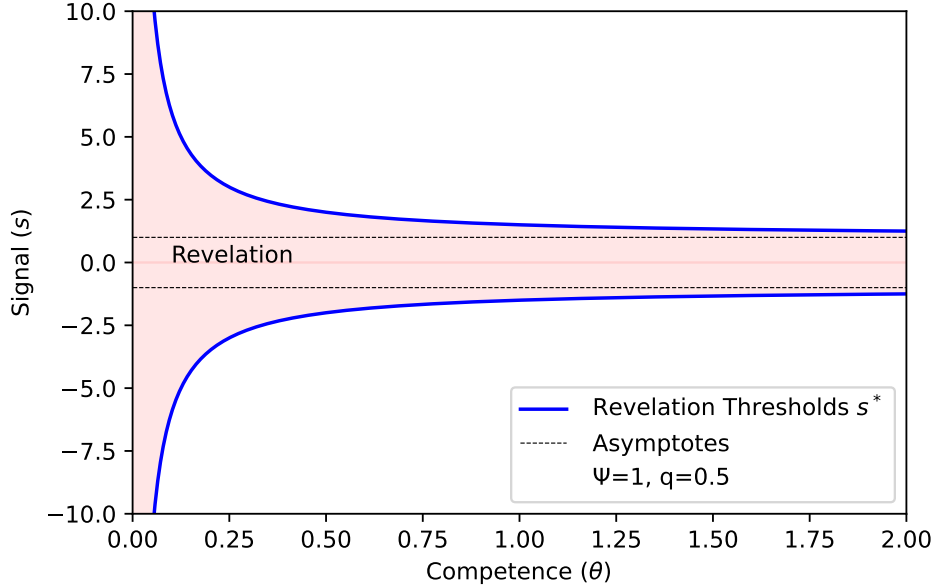


Figure 1: The solid lines show the signal thresholds $\hat{s}(\theta)$ as a function of the Advisor's qualification θ . When the Advisor receives a signal $s \geq [\hat{s}(\theta), \hat{s}(\theta)]$, she reveals her signal to the Leader. Dashed lines are asymptotes of the thresholds $\hat{s}(\theta)$ when $\Psi = 1$ and $q = 1/2$.

This specification of the actors' utilities allows us to give a precise description of the trade-off that the Leader faces between access to information and the quality of information. From Lemma 1, the probability that an Advisor of competence θ reveals her informative signal to the leader is $r(\theta) = Pr[s \geq (\hat{s}(\theta), \hat{s}(\theta))]$. The Leader's expected utility with an Advisor of competence θ is

$$EU_L[\theta] = r(\theta) \frac{1}{q + \theta} + (1 - r(\theta)) E[w^2 | s \notin (\hat{s}(\theta), \hat{s}(\theta))]. \quad (10)$$

It is important to note that, conditional on not observing an informative message from the Advisor, the Leader infers that the signal the Advisor observed is not in $(\hat{s}(\theta), \hat{s}(\theta))$. Therefore, the Leader's expected utility differs depending on the cause of the absence of an informative message. For the closed form of the Leader's expected utility see Appendix A.

The following proposition shows a key property of the Leader's expected utility as a function of agent types in the quadratic utilities model:


Proposition 3. *There exists a unique threshold $\Psi(q)$ such that the Leader’s equilibrium welfare is higher with an Advisor of some finite competence for all $\Psi < \Psi(q)$ but is highest with an Advisor of infinite competence otherwise.*

Proof. For proof see Appendix B. □

Proposition 3 shows that when the Advisors do not value office highly enough, the Leader is better off with an advisor of limited competence. This result may seem counter-intuitive, since the leader is always made better off by having better (more precise) information. However, advisors with high-quality information do not always deliver high-quality advice. Instead, as we show in Proposition 1, they are the most tempted to conceal their knowledge to avoid significant policy changes. The more highly the Advisor values the position, the less likely she is to conceal information. Because higher types conceal more information than do lower types, the Leader in a low Ψ setting is better off with an Advisor of limited competence who reveals her (low quality) information rather than with the more competent one who rarely provided (high quality) information.

Figure 2 shows the Leader’s expected utility as a function of competence (θ) for different Ψ . When the advisors value office highly (panel (a)), the Leader’s expected utility increases in the Advisor’s competence. However, as the office’s valuation decreases, the advisors of high competence begin to conceal more information from the Leader. Panel (b) illustrates the case in which the value of the position is low, and thus the Leader gets more useful advice (and higher expected utility) from a relatively low competence Advisor.

Figure 2: Expected Leader’s utility as a function of the Advisor’s competence for different Ψ



(a) The Leader’s expected utility when $\Psi = 10$

(b) The Leader’s expected utility when $\Psi = 0.5$

Comparative Statics

We can interpret Ψ as effectively measuring the Advisor’s opportunity cost of maintaining her position as the Advisor to the current Leader. Lowering Ψ means decreasing that opportunity

cost – i.e., making the outside options more attractive and the value of continuing as the Leader’s Advisor less attractive. We consider next how Ψ affects the Advisor’s optimal competence that maximizes the Leader’s utility, as well as, holding the Advisor’s competence constant, the Leader’s utility.

We can state the following result:

Proposition 4.

1. *The Leader’s equilibrium utility increases in the Advisor’s value of office, Ψ .*
2. *For $\Psi < \Psi(q)$, the type of Advisor θ that maximizes the Leader’s equilibrium utility is increasing in Ψ .*

Proof. See Appendix B for proof. □

We can interpret Ψ from an institutional perspective in a way that allows us to use our model to shed light on some of the under-appreciated effects of political polarization. Polarization decreases the correlation between preferences of political opponents and makes governance less common-value. It is reasonable to expect, thus, that higher polarization means that an Advisor who has been dismissed by a given Leader for what is, in effect, opposition to her policy agenda is more likely to find favor, at least in the short run, with the Leader’s political opponents. In effect, then, polarization decreases the value to the Advisor of keeping the current Leader satisfied, i.e., it lowers Ψ . With this interpretation in mind, Proposition 4, thus, suggests an underappreciated effect of polarization: With higher polarization, and so, lower Ψ , Leaders are induced to prefer lower-competence advisors. It also suggests that, as a consequence, they are less likely to choose policies that are radical departures from the status quo, since they are likely to remain relatively less informed.

Advisors with State-Dependent Preference

We show next that the basic incentives that drive the analysis above remain intact if we relax the assumption of the state-independence of the Advisor’s preferences. Suppose that both the Advisor’s and the Leader’s ideal policies depend on the state of the world, but the Advisor is more conservative than the Leader. The Leader’s bliss point is at w while the Advisor’s bliss point is at $c w$, where $c \in [0, 1]$ measures the Advisor’s conservatism, lower c corresponding to greater conservatism of the Advisor. When the Leader observes the signal, he sets policy to match the mean of the posterior distribution:

$$a = \frac{s \theta}{\theta + q}. \tag{11}$$

There exists an equilibrium, in which the Advisor reveals her information if and only if

$$\left(\frac{s}{\theta+q} - c \frac{s}{\theta+q}\right)^2 + \Psi > \left(c \frac{s}{\theta+q}\right)^2. \quad (12)$$

When the Advisor is not very conservative and shares the Leader's preference ($c > 1/2$), she reveals every signal s she observes. When the Advisor is sufficiently conservative, she reveals the signal she observes when it falls within $[\frac{\rho_{\bar{\Phi}}}{1+q/\theta} - \frac{\rho_{\bar{\Phi}}}{1+2c}, \frac{\rho_{\bar{\Phi}}}{1+q/\theta} + \frac{\rho_{\bar{\Phi}}}{1+2c}]$, and conceals all signals outside this interval. Note that the higher is the Advisor's conservatism (lower c), the wider is her revelation interval. This observation implies the following proposition:

Proposition 5. *When the Advisor's preference is state-dependent,*

1. *there exists a unique threshold $\Psi(q, c)$ such that the Leader's equilibrium welfare is higher with an Advisor of some finite competence for all $\Psi < \Psi(q, c)$ but is highest with an Advisor of infinite competence otherwise;*
2. *the optimal interior competence of the Advisor decreases in the Advisor's conservatism (i.e., increases in c);*
3. *the Leader's utility decreases in the Advisor's conservatism (i.e., increases in c).*

Proof. See Appendix G for the proof. □

Proposition 5 confirms that the baseline model's results are robust to a certain degree of common interest between the Advisor and the Leader. And, consistent with the results from prior studies, the more closely aligned their preferences, the more information the Leader receives in equilibrium.

Endogenous Information Acquisition

The analysis presented above proceeds in the setting where the Advisor's information about the state of the world is exogenously given, and her choice comes down to whether to share that information with the Leader. In this section, we consider a version of our model that incorporates endogenous information acquisition in the form of policy experiments – an especially relevant extension, given our interpretation of the Advisor as a bureaucratic agency.

Before proceeding, it is worth noting an important dimension of heterogeneity among settings in which the Advisor's choices on whether to run an experiment may take place: Conditional on an Advisor choosing to acquire information, some Advisors might lack the

ability to conceal information from the Leader, while others may be readily able to do so. Which of these possibilities is relevant might be driven by close connections between the relevant bureaucratic agency and the Leader, the particular importance of an issue, or the presence of whistle-blowers. We proceed by separately considering the implications of public and private experiments.

The game proceeds as follows. The Advisor of competence θ decides whether to conduct an experiment. If she does not, the Leader chooses a policy to implement, and the game ends. If she does, the Advisor chooses parameter $\tau \geq (0, \theta)$ that characterizes the precision of the experiment and observes the experiment realization: signal $s = w + \varepsilon$, where $\varepsilon \sim N(0, \tau)$. The Advisor, thus, can decide whether to “sabotage” the experiment, intentionally lowering the experiment’s precision to a level lower than her competence. We assume that the Leader does not observe τ chosen by the Advisor, but his conjecture $\hat{\tau}$ must be correct in any pure-strategy equilibrium, consistent with the definition of equilibrium. After that, depending on the nature of the experiment, the Leader either immediately observes the signal realization (if the experiment is public) or must await the Advisor’s decision to share this information (if the experiment is private). Once the revelation stage is complete, the Leader updates the policy based on the information he receives.

We begin with the following observation, which applies to both public and private experiment settings:

Remark 1. *Conditional on running the experiment, in equilibrium, the Advisor will choose the experiment of highest possible precision.*

Proof. See Appendix C. □

In effect, the Advisor can never credibly commit to sabotaging the experiment. Holding the Leader’s conjecture about τ , $\hat{\tau}$, fixed, increasing the precision of the experiment alters the signal realization distribution. The higher the precision, the less likely (in expectation) is the realization to fall further away from the Advisor’s ideal point, and therefore, the Advisor will always have an incentive to deviate from any precision level that is lower than the maximally value, θ . Therefore, $\hat{\tau} = \theta$ as the only value of τ that is consistent with equilibrium play is $\tau = \theta$, and so we proceed setting $\tau = \theta$ in the subsequent analysis.

Public Experiment. Consider now the public experiment environment, where the Leader immediately learns the signal the Advisor observes if the latter runs the experiment. In equilibrium, when the Advisor runs the experiment, she gets an office benefit Ψ and expects the Leader to update a policy to accommodate the experiment’s realization. When

the Advisor proceeds with the experiment ($R = 1$), she gets utility

$$EU_A(R = 1) = \int_{-1}^1 \left(\left(\frac{x}{1 + q/\theta} \right)^2 + \Psi \right) \phi\left(\frac{x}{\sqrt{1/q + 1/\theta}}\right) \sqrt{1/q + 1/\theta} dx, \quad (13)$$

where $\phi(x)$ is the probability density function of the standard normal distribution. When the Advisor chooses not run the experiment ($R = 0$), she gets utility zero. Because the decision to run the experiment inevitably leads to information revelation (because the experiment is public), the Advisor has to consider the Leader's reaction to every possible realization of the signal. Our next proposition summarizes the Advisor's equilibrium strategy and associated comparative statics.

Proposition 6. *When the experiment is public,*

1. *the Advisor runs the experiment when her competence is below the threshold $\hat{\theta}(q, \Psi)$ $\max\{0, \frac{q^2}{1+q}\}$ and does not run it otherwise; and*
2. *the threshold $\hat{\theta}(q, \Psi)$ increases in the Advisor's valuation of office Ψ and in the precision of the prior information q .*

Proof. See Appendix C. □

Note that the higher the Advisor's competence θ , the more the Leader updates the policy when he observes a signal realization. Therefore, the higher the Advisor's competence, the less likely she is to conduct the experiment. The higher the office benefit, the higher is the Advisor's utility from running the public experiment per se, and thus the more willing the Advisor of any competence is to run the experiment, and the higher is the competence threshold above which the Advisor never runs the experiment. The higher the precision of the prior (q), the less the Leader updates based on the information he receives, and the more likely the Advisor is to provide him with information by running the public experiment.

Private Experiment. Consider next the possibility of the private experiment. In this game, the Advisor first observes the signal realization and then decides whether to reveal this information to the Leader. The Advisor expects to receive

$$EU_A(R = 1) = \int_{s(\theta,)}^{s(\theta,)} \left(\left(\frac{x}{1 + q/\theta} \right)^2 + \Psi \right) \phi\left(\frac{x}{\sqrt{1/q + 1/\theta}}\right) \sqrt{1/q + 1/\theta} dx \quad (14)$$

when she decides to acquire information and utility zero if she never initiates the experiment. Equation 14 mirrors 13, but with the difference that, once the Advisor runs the experiment,

her only sequentially rational decision is to reveal her information when the signal she sees falls into the interval $[\hat{s}, \hat{s}]$ and to conceal information otherwise.

When the experiment is private, the Advisor’s control over the information flow allows her never to share the realizations that might lead to drastic policy changes. For this reason, when facing a choice of whether to run a private experiment, the Advisor always does so (assuming, of course, as we have been, that the experiment is costless to run).

Proposition 7. *When the experiment is private,*

1. *the Advisor always runs the experiment; and*
2. *the Advisor reveals the result of the experiment when it falls within the interval $[\hat{s}, \hat{s}]$ and conceals otherwise.*

Proof. See Appendix C. □

Proposition 6 and 7 explore how endogenous information acquisition affects the equilibrium availability of information to the Leader. What we discover in this section extends and reinforces the competence-revelation trade-off we study in the baseline model. Conditional on having acquired information, an Advisor of higher competence continues to have greater incentive to conceal this information, but, on top of this, this section shows that more competent Advisors are also (weakly) less likely to acquire information. The disincentive to acquire information is particularly acute when the information would necessarily be publicly available and less so when the information would be private.

Whistle-Blowing

In our model, the trade-off between the quality of advice and the likelihood of receiving can incline the Leader to prefer less competent advisors. However, other things being equal, the Leader trivially prefers more competent advice to less competent advice. In this section, we explore an institutional feature that might both counteract the Advisors’ incentives to conceal information from the Leader and improve the Leader’s information. This feature is the set of provisions that encourage whistle-blowing, i.e., the revelation of information outside the chain of command – in this case, from other agents who are privy to the advisor’s information (perhaps, the advisor’s subordinates within the bureaucracy) directly to the Leader. We show that when the Advisor’s information is fixed (exogenously given), the possibility of whistle-blowing (1) encourages revelation from an advisor of any competence and (2) improves the Leader’s utility premium from having a more competent Advisor. But when the Advisor’s

information is endogenous, whistle-blowing can backfire and lead to less information for the Leader.

To model whistle-blowing in a simple way, suppose that in the event of no revelation from the Advisor, the Leader gains access to the Advisor’s signal directly with probability ρ with the interpretation that a member of the bureaucratic apparatus to which the Advisor belongs bypasses her superior to reveal information to the Leader. While we are not modeling the motives of such a whistle-blower, parameter ρ flexibly captures the variability of the potential incentives: e.g., greater (lesser) protection for whistle-blowers, which, other things being equal, encourages (discourages) independent revelations. In what follows, then, we will, in a reduced-form way, refer to ρ as the probability of whistle-blowing.

In this setting, in equilibrium, the Advisor reveals her signal to the Leader when

$$\sqrt{\frac{\Psi}{1-\rho}} (1+q/\theta) < s < \sqrt{\frac{\Psi}{1-\rho}} (1+q/\theta) \hat{s}(\theta, \rho,). \quad (15)$$

Note that even, though the Advisor can choose the experiment’s precision, she always chooses $\tau = \theta$ in equilibrium. Therefore, the Advisor’s revelation strategy in equilibrium only depends on her true competence θ , not on the Leader’s conjecture ($\hat{\tau}$). For any level of competence θ , the Advisor reveals more information as ρ increases. Figure 5a shows thresholds that define the Advisor’s revelation for different levels of whistle-blowing. As ρ increases, the thresholds widen and the Advisor reveals more of the signals she may observe to the Leader.

Holding the Advisor’s information fixed, there are, thus, two channels by which whistle-blowing affects the Leader’s utility. First, the higher the value of ρ , the more likely the Leader is to receive information from the whistle-blower, other things being equal (the direct impact of whistle-blowing). Second, the higher the value of ρ , the greater are the Advisor’s incentives to reveal information (the indirect impact of whistle-blowing). Figure 5b depicts the Leader’s expected utility as a function of the Advisor’s competence given the Advisor’s information is exogenously fixed. Different curves represent the Leader’s expected utility without whistle-blowing; with the direct effect of whistle-blowing but suppressing its indirect effect; and with both the direct and the indirect effects of whistle-blowing together. Finally, because increasing the probability of whistle-blowing widens the Advisor’s revelation interval, the optimal interior competence of the Advisor is higher for higher ρ (see Appendix D for proof).

The following proposition summarizes these findings:

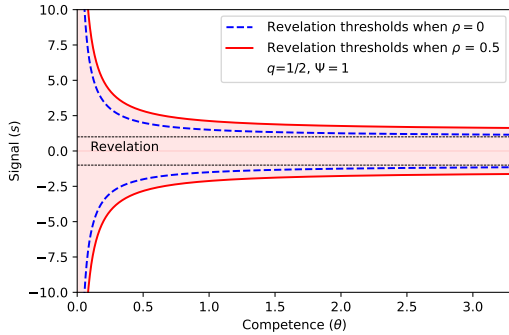
Proposition 8. *Given the Advisor’s information is exogenously fixed,*

1. *for any level of the Advisor’s competence θ , the higher the probability of whistle-blowing,*

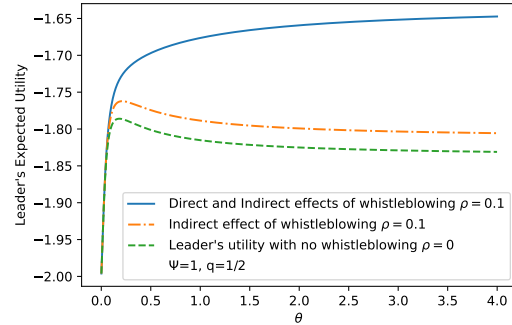
ρ , the more information the Advisor reveals to the Leader; and

- the Advisor’s interior competence θ that maximizes the Leader’s equilibrium welfare increases in the probability of whistle-blowing.

Figure 3: Whistle-blowing with exogenous information of the Advisor



(a) The Advisor’s revelation thresholds with whistle-blowing



(b) The Leader’s expected utility with whistle-blowing, given fixed information of the Advisor

Whistle-blowing can, thus, be understood as an institutional mechanism that allows the Leader to mitigate the trade-off between better vs. more information that we have described in the baseline model. Both of the changes described in Proposition 8 benefit the Leader, letting him get more informative advice and better adjust policies to the state of the world.

The relationships described in Proposition 8, however, depend importantly on the Advisor’s information being exogenously fixed. When the Advisor can choose whether to run an experiment, the probability of whistle-blowing alters the nature of this experiment. If the Advisor anticipates that the information she acquires may be revealed against her wish, it affects her incentives to seek this information. In what follows, then, we explore whether a higher probability of whistle-blowing always benefits the Leader when the Advisor has authority over whether to run the experiment.

We consider the private experiment environment, in which, if the Advisor decides to run the experiment, she observes the experiment’s realization and then decides whether to reveal this realization to the Leader. If the Advisor decides to conceal the signal she observes, then the whistle-blower reveals this information to the Leader with probability ρ . When probability ρ approaches zero, the experiment becomes effectively private: its results are only observed by the Advisor, and the Leader can acquire information only through direct revelation by the Advisor. When ρ approaches one, the Leader immediately learns all the information that the Advisor observes, and the experiment is effectively public.

Conditional on the probability of whistle-blowing ρ , the Advisor receives the following expected utility when she runs the experiment:

$$\begin{aligned}
EU_A(R = 1) &= \int_{\hat{s}(\theta, \rho)}^{s(\theta, \rho)} \left(\left(\frac{x}{1 + q/\theta} \right)^2 + \Psi \right) \phi\left(\frac{x}{\sqrt{1/q + 1/\theta}}\right) \sqrt{1/q + 1/\theta} dx \\
&\quad + \rho \left(\int_{\gamma}^{s(\theta, \rho)} \left(\frac{x}{1 + q/\theta} \right)^2 \phi\left(\frac{x}{\sqrt{1/q + 1/\theta}}\right) \sqrt{1/q + 1/\theta} dx \right. \\
&\quad \left. + \int_{\hat{s}(\theta, \rho)}^{+\infty} \left(\frac{x}{1 + q/\theta} \right)^2 \phi\left(\frac{x}{\sqrt{1/q + 1/\theta}}\right) \sqrt{1/q + 1/\theta} dx \right),
\end{aligned} \tag{16}$$

and she receives 0 when she decides against it.

The following proposition describes the effect of the possibility of whistle-blowing when the Advisor obtains information endogenously:

Proposition 9. *In the presence of the possibility of whistle-blowing,*

1. *there exists a unique threshold $\hat{\theta}(q, \Psi, \rho)$, such that the Advisor runs the experiment when her competence θ is below this threshold, and does not run the experiment otherwise;*
2. *the threshold $\hat{\theta}(q, \Psi, \rho)$ increases in Ψ and q and decreases in ρ .*

Proof. See Appendix C. □


The second part of Proposition 9 partially mirrors the comparative statics in Proposition 6. The higher the Advisor's valuation of office and the more precise the prior, the higher the Advisor's utility when she decides to run the experiment. Thus, the higher the competence threshold above which she chooses against running the experiment. In contrast, the higher the probability with which the experiment results are conveyed to the Leader against the Advisor's will, the lower her incentives to initiate the experiment, as higher ρ lowers her control over what happens in the aftermath of the experiment's realization. Figure 4a shows the Advisor's expected utility as a function of her competence if she decides to run the experiment. In this figure we can see that the higher is ρ , the lower is the Advisor's expected utility when she decides to proceed with the experiment.

Finally, the next proposition summarizes how the probability of whistle-blowing ρ affects the Leader's utility conditional on endogenous information acquisition by the Advisor.

Proposition 10.

1. *For any triple of the Advisor's competence $\theta > 0$, office valuation $\Psi > 0$, and precision of the prior $q > 0$, there exists a unique threshold $\rho(\theta, \Psi, q)$ such that the Leader's*

Figure 4: Whistle-blowing with endogenous information-acquisition by the Advisor



(a) The Advisor’s expected utility from running public experiment ($\rho = 0$), private experiment ($\rho = 0$), and the private experiment with a possibility of whistle-blowing ($\rho = 0.5$).

(b) The competence threshold below which the Advisor runs the experiment and above which the Advisor does not run the experiment as a function of the probability of whistle-blowing (ρ). The dashed vertical line represents a threshold to the left of which the Advisor of any competence always runs the experiment.

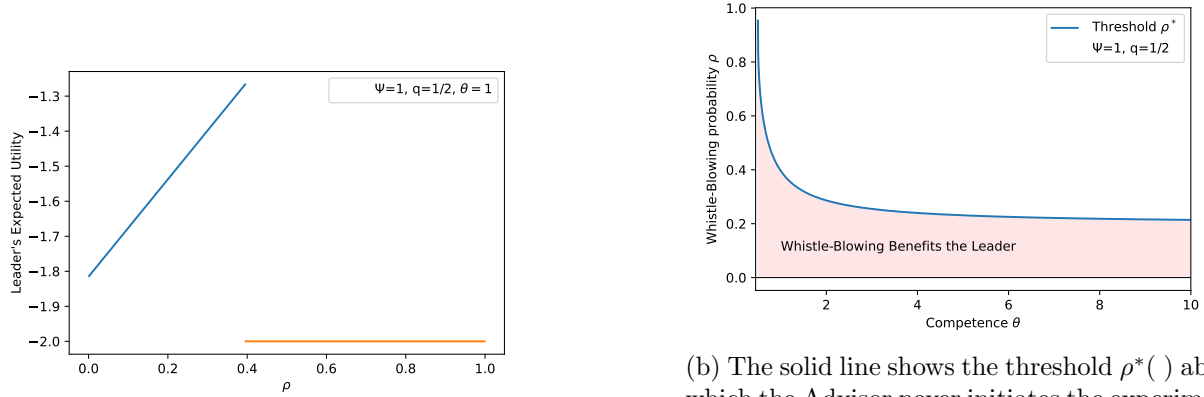
utility increases in the whistle-blowing probability ρ when $\rho < \rho^(\theta)$ and then drops to a constant value $K < EU_L(\rho = 0, \theta, \Psi, q)$.*

2. *The threshold $\rho^*(\theta, \Psi, q)$ decreases in θ .*

The higher the probability of whistle-blowing, the more likely is the Advisor to voluntarily reveal the information she observes, assuming the experiment is run. However, the higher is ρ , the lower is the threshold $\hat{\theta}(\rho)$ above which the Advisor refuses to run the experiment. Greater protection for whistle-blowers might, therefore, backfire and result in less information acquisition and worse outcomes for the Leader. Figure 5a depicts the Leader’s utility as a function of the whistle-blowing probability ρ for a given competence $\theta = 1$. As the whistle-blowing probability increases from 0, the Advisor, at first, reveals more information to the Leader. However, when the whistle-blowing probability ρ crosses the threshold $\rho^*(\theta = 1, \Psi = 1, q = 1/2)$, the Advisor prefers never to acquire information, as she anticipates that that information is too likely to be revealed to the Leader, against her will. Because of that, when ρ exceeds $\rho^*(\theta = 1, \Psi = 1, q = 1/2)$, the Leader observes no information about the state. The shaded area in Figure 5b shows the range of the whistle-blowing probability ρ that is beneficial to the leader. The solid curve describes the threshold ρ^* above which the Advisor never initiates the experiment, and, thus, above which the Leader never learns information about the state of the world.

The preceding analysis shows that the provisions offering protections to possible whistle-blowers are more desirable when the Advisor’s competence is lower and when Ψ and q are

Figure 5: Whistle-blowing with endogenous information acquisition by the Advisor: the effect on the Leader



(a) Two solid disjoint curves represent the Leader's expected utility as a function of the whistle-blowing probability ρ . When ρ is sufficiently high, the Advisor never initiates the experiment and the Leader's expected utility remains constant below $EU_L(\rho = 0,)$.

(b) The solid line shows the threshold $\rho^*(\theta)$ above which the Advisor never initiates the experiment. Below this curve, higher whistle-blowing probability always benefits the Leader as it encourages the Advisor to expand the revelation interval at $\Psi = 1, q = 1/2$, keeping the Advisor's competence fixed.

higher. But the Leader never wants whistle-blowing to become too likely (ρ too high) because that leaves the Leader with less information overall.

Discussion

TO BE ADDED

Appendices

A Non-monotonicity of the Leader's Preference

The Leader's expected utility:

$$E[U_L(\theta)] = \underbrace{Pr[s \in (\hat{s}, \hat{s})]}_{\text{Advisor sends informative message}} \underbrace{\frac{1}{q + \theta}}_{\text{Leader's expected utility after informative signal}} + \underbrace{Pr[s \notin (\hat{s}, \hat{s})]}_{\text{Advisor does not send informative message}} E[w^2 | s \notin (\hat{s}, \hat{s})], \quad (17)$$

where:

$$\begin{aligned}
A &= Pr[s \geq [\hat{s}, \hat{s}]] = \frac{1}{q + \theta} \\
&= (\Phi(\hat{s}/\sqrt{1/q + 1/\theta}) - \Phi(-\hat{s}/\sqrt{1/q + 1/\theta})) \frac{1}{q + \theta}.
\end{aligned} \tag{18}$$

and

$$\begin{aligned}
B &= Pr[s \notin (\hat{s}, \hat{s})] = E[w^2 | s \notin (\hat{s}, \hat{s})] \\
&= Pr[s < \hat{s}] E[w^2 | s < \hat{s}] + Pr[s > \hat{s}] E[w^2 | s > \hat{s}] \\
&= Pr[s < \hat{s}] \left(\int_{-\infty}^{\hat{s}} \int_{-\infty}^y x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s < \hat{s}]} dx dy \right) \\
&\quad + Pr[s > \hat{s}] \left(\int_{\hat{s}}^{\infty} \int_y^{\infty} x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s > \hat{s}]} dx dy \right) \\
&= \left(\int_{-\infty}^{\hat{s}} \int_{-\infty}^y x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/q}} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(\frac{x^2}{1/q} + \frac{y^2}{1/\theta})} dx dy \right) \\
&\quad + \left(\int_{\hat{s}}^{\infty} \int_y^{\infty} x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/q}} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(\frac{x^2}{1/q} + \frac{y^2}{1/\theta})} dx dy \right) \\
&= \frac{1}{2\pi} \frac{1}{\sqrt{1/q}} \frac{1}{\sqrt{1/\theta}} \int_{-\infty}^{\hat{s}} \frac{1}{q} \left(\frac{\rho_{\bar{q}}}{2\pi} e^{\frac{\theta y^2}{2}} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) \right) \right. \\
&\quad \left. + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) \right) + \frac{2}{\rho_{\bar{q}}} \frac{\Phi(\rho_{\bar{q}}(\hat{s} - y)) - \Phi(\rho_{\bar{q}}(\hat{s} + y))}{\rho_{\bar{q}}} dy \\
&= \int_{-\infty}^{\hat{s}} \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\theta y^2/2} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) \right. \\
&\quad \left. + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) \right) + \frac{2}{\rho_{\bar{q}}} \frac{\Phi(\rho_{\bar{q}}(\hat{s} - y)) - \Phi(\rho_{\bar{q}}(\hat{s} + y))}{\rho_{\bar{q}}} dy.
\end{aligned} \tag{19}$$

Let us denote

$$\begin{aligned}
g(a) &= \int_{-\infty}^a \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\theta y^2/2} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) \right. \\
&\quad \left. + \frac{2}{\rho_{\bar{q}}} \frac{\Phi(\rho_{\bar{q}}(\hat{s} - y)) - \Phi(\rho_{\bar{q}}(\hat{s} + y))}{\rho_{\bar{q}}} \right) dy.
\end{aligned} \tag{20}$$

Note that $g(1) = B$ and our objective is to compute $g(1)$. Let us start by computing $g(0)$

$$\begin{aligned}
g(0) &= \int_1^1 \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\theta y^2/2} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) + \frac{1}{\rho_{\bar{q}}} \right) dy \\
&= \frac{1}{q} \frac{e^{\frac{\hat{s}^2}{2(1/q+1/\theta)}} \sqrt{2/\pi} \hat{s}}{\sqrt{q(1+q/\theta)^3}}. \tag{21}
\end{aligned}$$

Now we compute derivative of $g(a)$ with respect to a . By Leibniz integral rule²

$$\begin{aligned}
\frac{\partial g(a)}{\partial a} &= \frac{\partial}{\partial a} \int_1^1 \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\theta y^2/2} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) \right. \\
&\quad \left. + \frac{2}{\rho_{\bar{q}}} \frac{\Phi(\rho_{\bar{a}q}(\hat{s} - y)) + \Phi(\rho_{\bar{a}q}(\hat{s} + y))}{\rho_{\bar{q}}} \right) dy \\
&= \int_1^1 \frac{\partial}{\partial a} \left(\frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\theta y^2/2} \left((\hat{s} - y) \phi(\rho_{\bar{q}}(\hat{s} - y)) + (\hat{s} + y) \phi(\rho_{\bar{q}}(\hat{s} + y)) \right. \right. \\
&\quad \left. \left. + \frac{2}{\rho_{\bar{q}}} \frac{\Phi(\rho_{\bar{a}q}(\hat{s} - y)) + \Phi(\rho_{\bar{a}q}(\hat{s} + y))}{\rho_{\bar{q}}} \right) \right) dy \\
&= \int_1^1 \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\frac{\theta y^2}{2}} \left(\phi(\rho_{\bar{a}q}(\hat{s} - y)) \frac{1}{2} \frac{\rho_{\bar{q}}}{a} (\hat{s} - y) - \phi(\rho_{\bar{a}q}(\hat{s} + y)) \frac{1}{2} \frac{\rho_{\bar{q}}}{a} (\hat{s} + y) \right) dy \\
&= \int_1^1 \frac{1}{2\pi} \frac{1}{\sqrt{q/\theta}} e^{\frac{\theta y^2}{2}} \frac{1}{2} \frac{\rho_{\bar{q}}}{a} \left(e^{-\frac{aq(\hat{s}-y)^2}{2}} (\hat{s} - y) + e^{-\frac{aq(\hat{s}+y)^2}{2}} (\hat{s} + y) \right) dy \\
&= \int_1^1 e^{-\frac{\theta y^2}{2}} \frac{\rho_{\bar{q}}}{\theta} \frac{e^{-\frac{aq(\hat{s}-y)^2}{2}} (\hat{s} - y) + e^{-\frac{aq(\hat{s}+y)^2}{2}} (\hat{s} + y)}{4\pi \rho_{\bar{a}q}} dy \\
&= \hat{s} \frac{e^{-\frac{aq\hat{s}^2\theta}{2(1+\frac{aq}{\theta})}}}{2\pi \rho_{\bar{a}q} (1 + \frac{aq}{\theta})^{3/2}}. \tag{22}
\end{aligned}$$

We now take an integral wrt a of $\frac{dg(a)}{da}$:

$$\begin{aligned}
g(a) &= \int \hat{s} \frac{e^{-\frac{aq\hat{s}^2\theta}{2(aq+\theta)}}}{2\pi \rho_{\bar{a}q} (1 + \frac{aq}{\theta})^{3/2}} da \\
&= \frac{2\Phi\left(\frac{\hat{s}}{\sqrt{\frac{1}{aq} + \frac{1}{\theta}}}\right)}{q} + C, \tag{23}
\end{aligned}$$

²Leibniz integral rule applies because the integral of partial derivative converges [link](#)

where C is unknown constant. Finally, let us note that

$$g(a = 0) = \frac{1}{q} \frac{e^{\frac{qs^2\theta}{2(q+\theta)}} \sqrt{2/\pi\hat{s}}}{\mathcal{P}\frac{1}{q}(1+q/\theta)^{3/2}} \quad (24)$$

and

$$g(a = 0) = \lim_{a \downarrow 0} \frac{2\Phi\left(\frac{\hat{s}}{\sqrt{\frac{1}{aq} + \frac{1}{\theta}}}\right) - 1}{q} + C = C, \quad (25)$$

where equation 24 is corollary of equation 21 and equation 25 is corollary of equation 23.

Therefore

$$C = \frac{1}{q} \frac{e^{\frac{qs^2\theta}{2(q+\theta)}} \sqrt{2/\pi\hat{s}}}{\mathcal{P}\frac{1}{q}(1+q/\theta)^{3/2}}. \quad (26)$$

Finally,

$$\begin{aligned} B &= \frac{\Pr[s \notin [\hat{s}, \hat{s}]]}{\Pr[s < \hat{s}]} g(a = 1) \\ &= \lim_{a \downarrow 1} \frac{2\Phi\left(\frac{\hat{s}}{\sqrt{\frac{1}{aq} + \frac{1}{\theta}}}\right) - 1}{q} \frac{1}{q} \frac{e^{\frac{qs^2\theta}{2(q+\theta)}} \sqrt{2/\pi\hat{s}}}{\mathcal{P}\frac{1}{q}(1+q/\theta)^{3/2}} \\ &= \left(\frac{2\Phi\left(\frac{\hat{s}}{\sqrt{\frac{1}{q} + \frac{1}{\theta}}}\right) - 1}{q} \frac{1}{q} \frac{e^{\frac{qs^2\theta}{2(q+\theta)}} \sqrt{2/\pi\hat{s}}}{\mathcal{P}\frac{1}{q}(1+q/\theta)^{3/2}} \right) \end{aligned} \quad (27)$$

Therefore, the expected Leader's utility is

$$\begin{aligned} E[U_L(\theta)] &= \left(2\Phi\left(\frac{\hat{s}}{\sqrt{1/q + 1/\theta}}\right) - 1 \right) \frac{1}{q + \theta} \\ &+ \left(\frac{2\Phi\left(\frac{\hat{s}}{\sqrt{\frac{1}{q} + \frac{1}{\theta}}}\right) - 1}{q} \frac{1}{q} \frac{e^{\frac{qs^2\theta}{2(q+\theta)}} \sqrt{2/\pi\hat{s}}}{\mathcal{P}\frac{1}{q}(1+q/\theta)^{3/2}} \right). \end{aligned} \quad (28)$$

A.1 Conservatism

This proposition is an immediate implication of a comparative statics analysis wrt Ψ . To see that, denote $\frac{1}{1-2c}$ as new Ψ .

B Existence of the finite optimum

The Leader's utility derivative wrt θ is:

$$\frac{\partial E[U_L(\theta)]}{\partial \theta} = \frac{1}{2(q+\theta)^2} \left(\frac{\rho}{2} e^{-\frac{\Psi q(1+q/\theta)}{2}} \frac{\sqrt{\frac{q\theta(q+\theta)}{\pi}} (2\theta + \Psi q(q+\theta))}{\theta^2} + \underbrace{2(2\Phi(\frac{\hat{s}}{\sqrt{1/q + 1/\theta}}) - 1)}_{>0} \right). \quad (29)$$

Note that $\frac{\partial E[U_L(\theta)]}{\partial \theta}$ converges to $\frac{1}{q^2}$ as θ approaches 0 and converges to 0 as θ approaches infinity, where the sign depends on whether

$$F(\Psi, q) = 4 - e^{-\frac{\Psi q}{2}} \sqrt{\frac{2}{\pi}} \sqrt{\Psi q} (2 + \Psi q) - 4\Phi(\sqrt{\Psi q})$$

is positive or negative. $F(\Psi = 0, q) = 0$ and it decreases in Ψ for $\Psi < 1/q$ and increases in Ψ when $\Psi > 1/q$. Therefore, for $\Psi \in (0, 1/q)$, $\frac{\partial E[U_L(\theta)]}{\partial \theta}$ is negative at $\theta = 1$ and, because $E[U_L(\theta)]$ is continuous and differentiable, the interior optimum will exist (note that $\Psi < 1/q$ is sufficient but not necessary condition for the existence of the interior optimum).

Let us denote the derivative of the Leader's utility with respect to the Advisor's qualification θ as $D(\theta, \Psi) = \frac{\partial E[U_L(\theta)]}{\partial \theta}$. Note that the Leader's utility reaches local maximum at $\hat{\theta}$ s.t. $D(\hat{\theta}, \Psi) = 0$. By the implicit function theorem, in order to seek how Ψ affects the Leader's choice of interior optimal Advisor's qualification, one needs to compute

$$\frac{d\hat{\theta}(\Psi)}{d\Psi} = \frac{\partial D(\theta, \Psi)}{\partial \theta D(\theta, \Psi)}. \quad (30)$$

Because we are looking for θ that maximizes the Leader's utility, $\partial_\theta D(\theta, \Psi)$ should not exceed zero. Therefore, sign of equation 30 mirrors sign of $\partial_\theta D(\theta, \Psi)$. Because

$$\partial_\theta D(\theta, \Psi) = \frac{e^{-\frac{\Psi q\theta(q+\theta)}{2\theta}}}{2} \frac{\Psi}{\theta^2} \underbrace{\frac{\rho}{2\pi} \frac{q^2}{\sqrt{\Psi q\theta(q+\theta)}}}_{>0} (\Psi q(q+\theta) - \theta), \quad (31)$$

sign of $(\Psi q(q+\theta) - \theta)$ determines whether $\hat{\theta}$ increases or decreases in Ψ . When $(\Psi q(q+\theta) - \theta)$ is positive, optimal interior competence increases in Ψ , and it decreases in Ψ when $(\Psi q(q+\theta) - \theta)$ is negative.

Note that $D(\theta, \Psi)$ reaches minimum at $\Psi = \frac{\theta}{q(q+\theta)}$. Next, because $D(\theta = 0, \Psi) = \frac{\rho}{2q^2}$

is positive while $D(\theta = \frac{q^2}{1}, \Psi) = \frac{(1-q)^2(3\frac{\rho_-}{2} + \frac{\rho_- e\pi 2(1-2)}{2^2 e\pi q^2})}{2^2 e\pi q^2}$ is negative, when $\Psi > \frac{\theta}{q(q+\theta)}$, there will be interior maximum of the Leader's utility $\hat{\theta}$ s.t. $\hat{\theta} < \frac{q^2}{1}$. Because $\Psi > \frac{\theta}{q(q+\theta)}$, this interior maximum ($\hat{\theta}$) increases in Ψ .

Finally, we prove that if the expected Leader's utility has a local maximum, this local maximum is unique. Once we prove this statement, we prove that **any** interior maximum increases in Ψ . First, note that derivative of the Leader's utility wrt θ is:

$$\frac{\partial E[U_L(\theta)]}{\partial \theta} = \underbrace{\frac{1}{2(q+\theta)^2}}_{>0} \left(\frac{\rho_-}{2} e^{-\frac{\Psi q(1+q/\theta)}{2}} \sqrt{\frac{q\theta(q+\theta)}{\pi}} (2\theta + \Psi q(q+\theta))}{\theta^2} + 2(2\Phi(\frac{\hat{s}}{\sqrt{1/q + 1/\theta}}) - 1) \right). \quad (32)$$

Sign of $\frac{\partial E[U_L(\theta)]}{\partial \theta}$ mirrors the sign of

$$Interior(\theta) = \frac{\rho_-}{2} e^{-\frac{\Psi q(1+q/\theta)}{2}} \sqrt{\frac{q\theta(q+\theta)}{\pi}} (2\theta + \Psi q(q+\theta))}{\theta^2} + 2(2\Phi(\frac{\hat{s}}{\sqrt{1/q + 1/\theta}}) - 1). \quad (33)$$

Now note that

$$\frac{\partial Interior(\theta)}{\partial \theta} = \frac{e^{-\frac{\Psi q(1+q/\theta)}{2}} q \sqrt{\frac{q\theta(q+\theta)}{2\pi}} \Psi (q\theta - \Psi q^2(q+\theta))}{\theta^4}. \quad (33)$$

Therefore, $Interior(\theta)$ decreases in θ for $\theta < \frac{q^2}{1}$ and increases in θ for $\theta > \frac{q^2}{1}$. It implies that the expected Leader's utility can have no more than one local maximum.

As we just proved, the interior maximum of the Leader's utility exists if and only if $\Psi > \frac{\theta}{q(q+\theta)}$. Therefore, when the interior maximum exists,

$$\partial D(\theta, \Psi) = \frac{e^{-\frac{\Psi q\theta(q+\theta)}{2\theta}} \Psi q^2}{2 \theta^2 \rho_- \frac{2\pi}{\sqrt{\Psi q \theta (q+\theta)}}} (\Psi q(q+\theta) - \theta) > 0.$$

Therefore, for any q there exists a unique threshold $\Psi(q)$ s.t. the Leader prefers interior competence over infinite competence for every $\Psi < \Psi(q)$.

C Experiment

C.1 Credible Precision

Assume that the Leader's beliefs of the precision the Advisor chooses are fixed at $\hat{\tau}$. Then, the Advisor's utility from running the private experiment is

$$\begin{aligned} & Pr[s \geq (\frac{\rho - \Psi}{1 + \frac{q}{\hat{\tau}}}, \frac{\rho - \Psi}{1 + \frac{q}{\hat{\tau}}})] E[U_A(R = 1)] \\ &= (\Phi(\frac{\rho - \Psi}{1/q + 1/\tau}) - \Phi(\frac{\rho - \Psi}{1/q + 1/\hat{\tau}})) \\ & \int_{\frac{\rho - \Psi}{1/q + 1/\hat{\tau}}}^{\frac{\rho - \Psi}{1/q + 1/\tau}} \left(\left(\frac{x}{1 + q/\hat{\tau}} \right)^2 + \Psi \right) \phi\left(\frac{x}{\sqrt{1/q + 1/\tau}}\right) \sqrt{1/q + 1/\tau} dx, \end{aligned} \quad (34)$$

and the Advisor's utility from running the public experiment is

$$\int_{\frac{\rho - \Psi}{1/q + 1/\tau}}^{\frac{\rho - \Psi}{1/q + 1/\tau}} \left(\left(\frac{x}{1 + q/\hat{\tau}} \right)^2 + \Psi \right) \phi\left(\frac{x}{\sqrt{1/q + 1/\tau}}\right) \sqrt{1/q + 1/\tau} dx. \quad (35)$$

Both increase in τ .

The Advisor's utility conditional on revelation always increase in τ . The higher precision experiment is less likely to generate an outcome further away from the most preferred Advisor's outcome. The Advisor's probability of receiving information that she will reveal to the Leader also increases in τ because the higher precision experiment is less likely to generate an outcome further away from the revelation interval.

C.2 Public Experiment

When the experiment is public and the Advisor decides to run it, she gets utility

$$\Psi - \frac{1}{q(1 + q/\theta)}. \quad (36)$$

It exceeds zero when $\theta < \hat{\theta} = \max\{1, \frac{q^2}{1 - \Psi}\}$, 0 and is less than zero otherwise.

The threshold $\hat{\theta}$ weakly increases in Ψ because

$$\frac{\partial \hat{\theta}}{\partial \Psi} = \frac{q^2}{(1 - \Psi - q)^2} > 0, \quad (37)$$

and increases in q because

$$\frac{\partial \hat{\theta}}{\partial q} = \frac{\Psi - q(2 - \Psi - q)}{(1 - \Psi - q)^2}, \quad (38)$$

which is positive when $q < 2/\Psi$ and negative otherwise, but when $q > 2/\Psi$, the threshold $\hat{\theta}$ is zero. Therefore, $\hat{\theta}$ weakly increases in q .

C.3 Private Experiment

When the experiment is private, the advisor gets utility

$$E_A(R = 1) = \frac{\phi\left(\frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}}\right) \frac{2}{\Psi} + \frac{\left(\frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}}\right) \frac{1}{q/\theta} \left(\frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}}\right) \frac{1}{(q+\theta)}}{\rho \frac{s}{1/q+1/\theta} \sqrt{1+q/\theta}}. \quad (39)$$

Its derivative wrt θ is

$$\frac{\partial E_A(R = 1)}{\partial \theta} = \frac{2\phi\left(\frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}}\right) \frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}} \left(2\Phi\left(\frac{\rho \frac{s}{1/q+1/\theta}}{\rho \frac{s}{1/q+1/\theta}}\right) - 1\right) \left(1 - \rho\right)}{q^2(q+\theta)^2}, \quad (40)$$

and the derivative of 40 wrt Ψ is equal to

$$\frac{1}{q^2(q+\theta)^2} \frac{1}{\Psi} \phi\left(\frac{s}{\sqrt{1/q+1/\theta}}\right) \rho \frac{\rho}{\pi} \left(\frac{s}{\sqrt{1/q+1/\theta}}\right)^3 < 0. \quad (41)$$

Therefore, $\frac{\partial E_A(R=1)}{\partial \theta}$ decreases in Ψ and reaches maximum when Ψ converges to zero. Thus, the maximum possible $\frac{\partial E_A(R=1)}{\partial \theta}$ is

$$\lim_{\Psi \rightarrow 0} \frac{\partial E_A(R = 1)}{\partial \theta} = 0. \quad (42)$$

Thus, $E_A(R = 1)$ decreases in θ and reaches minimum when θ converges to infinity, where

$$\lim_{\theta \rightarrow +\infty} E_A(R = 1) = 0. \quad (43)$$

Then the Advisor always runs the private experiment.

D Whistle-Blowers

When the Advisor reveals her signal (s), she gets utility

$$\left(\frac{s}{1+q/\theta}\right)^2 + \Psi; \quad (44)$$

and when she conceals her information, she gets utility

$$\rho \left(\frac{s}{1 + q/\theta} \right)^2 - (1 - \rho) d(\theta). \quad (45)$$

Therefore, she reveals her signal when

$$\sqrt{d(\theta) + \frac{\Psi}{1 - \rho}} (1 + q/\theta) < s < \sqrt{d(\theta) + \frac{\Psi}{1 - \rho}} (1 + q/\theta) \quad (46)$$

and conceals otherwise. Because the revelation interval is symmetric around zero, $d(\theta) = 0$. Thus, the Advisor reveals her signal if and only if

$$\sqrt{\frac{\Psi}{1 - \rho}} (1 + q/\theta) < s < \sqrt{\frac{\Psi}{1 - \rho}} (1 + q/\theta) > \hat{s}. \quad (47)$$

Now note that the direct effect of whistleblowing always incentivizes the Leader to favor more competent advisors, regardless of which competence he preferred with no whistleblowing present. When it comes to the indirect effect of whistleblowing, the higher ρ is equivalent to having an advisor who values office as $\frac{1}{1 - \rho}$ instead of Ψ . As we know, the optimal Advisor's competence increases in Ψ . Therefore, the optimal Advisor's competence will increase in ρ .

D.1 Endogenous Information acquisition

When the Advisor decides to run the experiment, the derivative of her utility wrt θ is

$$\begin{aligned} & \frac{\partial EU_A(run)}{\partial \theta} \\ = & \frac{\rho + 2\phi\left(\sqrt{\frac{q(1+q/\theta)}{1 - \rho}}\right) \sqrt{\Psi q (1 + q/\theta)} (1 - \rho) - (2\Phi\left(\sqrt{\frac{q(1+q/\theta)}{1 - \rho}}\right) - 1) (1 - \rho)}{q^2(q + \theta)^2}, \end{aligned} \quad (48)$$

which decreases in Ψ

$$\begin{aligned} & \frac{\partial^2 EU_A(run)}{\partial \theta \partial \Psi} \\ = & \phi\left(\sqrt{\frac{\Psi q(1 + q/\theta)}{2(1 - \rho)}}\right) \sqrt{\frac{\Psi q}{(1 - \rho)(1 + q/\theta)}} \frac{1}{\theta} < 0. \end{aligned} \quad (49)$$

Therefore, $\frac{\partial EU_A(run)}{\partial \theta}$ reaches maximum as Ψ approaches zero. Now note that as Ψ approaches 0, $\frac{\partial EU_A(run)}{\partial \theta}$ converges to $\frac{\rho}{\pi q^2} < 0$. Therefore, the Advisor's utility from running the experiment always decreases in θ and the higher is the Advisor's competence, the lower are her incentives to initiate an experiment.

By the implicit function theorem, $\frac{\partial \theta(\cdot)}{\partial \rho} = \frac{\partial_{\Psi} EU_A(run)}{\partial_{\theta} EU_A(run)}$ the denominator of this fraction is always negative. The numerator of this fraction is equal to

$$\frac{\partial EU_A(run)}{\Psi} = 2\Phi\left(\sqrt{\frac{\Psi}{1-\rho}} \frac{q(1+q/\theta)}{1-\rho}\right) - 1 < 0. \quad (50)$$

Therefore, θ increases in Ψ . Similarly, because the expected Advisor's utility from running the experiment increases in q , the threshold θ increases in q .

Finally, note that the Advisor's utility from running an experiment is decreasing in ρ . The derivative of the Advisor's utility wrt ρ is

$$\begin{aligned} & \frac{\partial EU_A(run)}{\partial \rho} \\ &= \frac{1}{q(1+q/\theta)} + \frac{2\Phi\left(\sqrt{\frac{q(1+q/\theta)}{1-\rho}}\right) - 1}{q(1+q/\theta)} - \frac{e^{-\frac{\Psi q(1+q/\theta)}{2(1-\rho)}}}{\sqrt{2\pi(1-\rho)}}, \end{aligned} \quad (51)$$

which is increasing in Ψ

$$\begin{aligned} & \frac{\partial^2 EU_A(run)}{\partial \rho \partial \Psi} \\ &= \frac{e^{-\frac{\Psi q(1+q/\theta)}{2(1-\rho)}}}{\sqrt{2\pi(1-\rho)}} \frac{\sqrt{\Psi q(1+q/\theta)}}{(1-\rho)} > 0. \end{aligned} \quad (52)$$

Therefore, $\frac{\partial EU_A(run)}{\partial \rho}$ reaches maximum as Ψ approaches ∞ . As Ψ converges to ∞ , $\frac{\partial EU_A(run)}{\partial \rho}$ converges to zero. Thus, the threshold θ below which the Advisor reveals her experiment decreases in ρ .

References

- [1] Kenneth W Abbott et al. "Competence versus control: The governor's dilemma". In: *Regulation & Governance* 14.4 (2020), pp. 619–636.
- [2] David Austen-Smith. "Information transmission in debate". In: *American Journal of political science* (1990), pp. 124–152.

- [3] David Austen-Smith. “Interested experts and policy advice: Multiple referrals under open rule”. In: *Games and Economic Behavior* 5.1 (1993), pp. 3–43.
- [4] Ying Bai and Titi Zhou. ““Mao’s last revolution”: a dictator’s loyalty–competence tradeoff”. In: *Public Choice* 180.3 (2019), pp. 469–500.
- [5] Marco Battaglini. “Multiple referrals and multidimensional cheap talk”. In: *Econometrica* 70.4 (2002), pp. 1379–1401.
- [6] Sourav Bhattacharya, Maria Goltsman, and Arijit Mukherjee. “On the optimality of diverse expert panels in persuasion games”. In: *Games and Economic Behavior* 107 (2018), pp. 345–363.
- [7] Sourav Bhattacharya and Arijit Mukherjee. “Strategic information revelation when experts compete to influence”. In: *The RAND Journal of Economics* 44.3 (2013), pp. 522–544.
- [8] Yeon-Koo Che and Navin Kartik. “Opinions as incentives”. In: *Journal of Political Economy* 117.5 (2009), pp. 815–860.
- [9] Wioletta Dziuda. “Strategic argumentation”. In: *Journal of Economic Theory* 146.4 (2011), pp. 1362–1397.
- [10] Georgy Egorov and Konstantin Sonin. “Dictators and their viziers: Endogenizing the loyalty–competence trade-off”. In: *Journal of the European Economic Association* 9.5 (2011), pp. 903–930.
- [11] Sean Gailmard and John W Patty. “Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise”. In: *American Journal of Political Science* 51.4 (2007), pp. 873–889.
- [12] Thomas W Gilligan and Keith Krehbiel. “Asymmetric information and legislative rules with a heterogeneous committee”. In: *American journal of political science* (1989), pp. 459–490.
- [13] Aziz Huq and Tom Ginsburg. “How to lose a constitutional democracy”. In: *UCLA L. Rev.* 65 (2018), p. 78.
- [14] Ora John Reuter and Graeme B Robertson. “Subnational appointments in authoritarian regimes: Evidence from Russian gubernatorial appointments”. In: *The Journal of politics* 74.4 (2012), pp. 1023–1037.
- [15] Victor Shih, Christopher Adolph, and Mingxing Liu. “Getting ahead in the communist party: explaining the advancement of central committee members in China”. In: *American political science review* 106.1 (2012), pp. 166–187.

- [16] Hyun Song Shin. “The burden of proof in a game of persuasion”. In: *Journal of Economic Theory* 64.1 (1994), pp. 253–264.
- [17] Joel Sobel. “Giving and receiving advice”. In: *Advances in economics and econometrics* 1 (2013), pp. 305–341.
- [18] Joel Sobel. “Information control in the principal-agent problem”. In: *International Economic Review* (1993), pp. 259–269.
- [19] Kimiko Terai and Amihai Glazer. “Rivalry among agents seeking large budgets”. In: *Journal of Theoretical Politics* 30.4 (2018), pp. 388–409.
- [20] Asher Wolinsky. “Information transmission when the sender’s preferences are uncertain”. In: *Games and Economic Behavior* 42.2 (2003), pp. 319–326.
- [21] Tianyang Xi. “All the emperor’s men? Conflicts and power-sharing in imperial China”. In: *Comparative Political Studies* 52.8 (2019), pp. 1099–1130.